

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:)
Douglas M. Dillon, et al.) : Examiner: Nghi V. Tran
Appln. No.: 10/010,521) : Group Art Unit: 2151
Filed: December 7, 2001) : Confirmation No.: 1352
For: METHOD AND APPARATUS FOR)
SELECTIVELY ALLOCATING)
AND ENFORCING BANDWIDTH)
USAGE REQUIREMENTS ON)
NETWORK USERS :
:

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

DECLARATION UNDER 37 C.F.R. § 1.131 OF
DOUGLAS M. DILLON AND VIVEK GUPTA

Sir:

We, Douglas M. Dillon and Vivek Gupta, hereby declare that:

1. We are the inventors of the inventions described and claimed in the above-identified patent application.

Claims 59 through 61

2. Prior to December 23, 2007, we conceived in the United States of the inventions set forth in Claims 59 through 61 of the above-identified patent application, as set forth in the Claim Sheet attached as Exhibit 1.

3. Also prior to that date, we prepared documents entitled "DirecPC Turbo-Throttling Feature Design Document" and "Turbo-Throttling II Data Sheet".

4. Copies of the documents, from which dates have been redacted, are attached as Exhibits 2 and 3.

5. These documents evidence the conception of the inventions of the claims.

6. In particular, with respect to Claims 59 through 61, conception is evidenced, e.g., by the references in the documents to the hybrid gateway, to maintenance of average data rate or running average throughput using a leaky bucket, wherein if a user goes beyond the average data rate, he will be throttled through a reduction in the TCP window size, and to per-IP address statistics including TIServiceID (service plan to which the user belongs). (See, e.g., Exhibit 2, pp. 9-10; Exhibit 3.)

7. Thereafter, during the time period from before December 23, 1997, through February 6, 1998, we worked on designing and constructing an actual device implementing the inventions of the foregoing claims.

8. For example, in January 1998, we revised our design specification to include additional details regarding throttling.

9. By February 6, 1998, we had constructed the actual device in the U.S.

10. The device operated successfully by that date.

11. This is evidenced by an email that my coworker Harvey Lindenbaum sent on February 9, 1998.

12. In that email, which was a weekly report for the week ending on February 6, 1998, he indicated that "FIT testing of Turbo Throttling 2 has been completed".

13. This means that all aspects of the project were successfully tested.

14. A copy of his email, from which information relating to matters other than Turbo Throttling has been redacted, is attached as Exhibit 4.

Claims 23, 33, 53, and 56

15. Furthermore, prior to April 28, 2008, we conceived in the United States of the inventions set forth in Claims 23, 33, 53, and 56 of the above-identified patent application, as set forth in the Claim Sheet attached as Exhibit 1.

16. Also prior to that date, we prepared a description entitled "DirecPC Turbo-Throttling Feature Requirements Specification".

17. A copy of the document, from which dates have been redacted, is attached as Exhibit 5.

18. The document evidences the conception of the inventions of the claims.

19. With respect to Claims 23, 33, 53, and 56, conception is evidenced, e.g., by the references in the document to the hybrid gateway (HGW) utilizing a "service plan" and effecting throttling "per user", by the references in the document to the number of TCP connections being limited to the MaxConnections configured for the corresponding service plan, and by the references in the document to "thresholds". (See, e.g., Exhibit 5, §§ 2.1 and 3.1.1-3.1.7 and "Turbo-Throttling II Data Sheet".)

20. Prior to April 28, 2008, we constructed and successfully tested in the U.S. an actual device implementing the inventions of the foregoing claims, as explained above in paragraphs 9 through 14.

21. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of this application and any patent issuing thereon.

6/17/2008
Date (Month/Day/Year)

7/1/2008
Date (Month/Day/Year)

Douglas M. Dillon
Douglas M. Dillon

Vivek Gupta
Vivek Gupta

EXHIBIT 1

Appln. No. 10/010,521 Claims 23, 33, 53, 56, and 59 through 61

23. A gateway for use in a system wherein a first apparatus, said gateway, and a second apparatus are in a TCP/IP network, each of the first apparatus, said gateway, and the second apparatus having different IP addresses, said gateway comprising:

a throttling unit that is configured to (a) determine the number of TCP connections that are open and (b) throttle a user of the first apparatus in accordance with (1) the determination of the number of TCP connections that are open and (2) a level of service subscribed to by the user of the first apparatus.

33. A method comprising:

determining a number of TCP connections that are open; and
throttling, by a gateway for use in a system wherein a first apparatus, the gateway, and a second apparatus are in a TCP/IP network, of a user of the first apparatus, in accordance with (1) the determination of the number of TCP connections that are open and (2) a level of service subscribed to by the user.

53. A gateway for use in a system wherein a first apparatus, said gateway, and a second apparatus are in a TCP/IP network, each of the first apparatus, said gateway, and the second apparatus having a different IP address, said gateway comprising:

throttling means for determining a number of TCP connections that are open and for throttling a user of the first apparatus, in accordance with (1) the determination of the number of TCP connections that are open and (2) a level of service subscribed to by the user.

56. An apparatus according to Claim 53, wherein said throttling means compares bandwidth usage to a threshold.

59. A gateway for use in a system wherein a first apparatus, said gateway, and a second apparatus are in a TCP/IP network, each of the first apparatus, said gateway, and the second apparatus having different IP addresses, said gateway comprising:

a determining unit that is configured to determine which of a plurality of service plans a user of the first apparatus subscribes to; and

a throttling unit that is configured to throttle the user in accordance with (1) a leaky bucket analysis of the user's throughput and (2) the service plan subscribed to by the user as determined by said determining unit,

wherein said throttling unit intercepts a packet on a TCP/IP connection between the first apparatus and the second apparatus; and

wherein said throttling unit effects throttling by modifying a field in the packet to cause the second apparatus to change an amount of data it sends before awaiting a TCP ACK from the first apparatus.

60. A method comprising:

determining by a gateway, for use in a system wherein a first apparatus, the gateway, and a second apparatus are in a TCP/IP network, which of a plurality of service plans a user of the first apparatus subscribes to;

throttling by the gateway of a user of the first apparatus, in accordance with (1) a leaky bucket analysis of the user's throughput and (2) the service plan subscribed to by the user as determined by said determining step,

wherein the first apparatus, the gateway, and the second apparatus have different IP addresses,

wherein the gateway intercepts a packet on a TCP/IP connection between the first apparatus and the second apparatus, and

wherein said throttling comprises modifying a field in the packet to cause the second apparatus to change an amount of data it sends before awaiting a TCP ACK from the first apparatus.

61. A gateway for use in a system wherein a first apparatus, said gateway, and a second apparatus are in a TCP/IP network, said gateway comprising:

determining means for determining which of a plurality of service plans a user of the first apparatus subscribes to; and

throttling means for throttling the user, in accordance with (1) a leaky bucket analysis of a user's throughput and (2) the service plan subscribed to by the user as determined by said determining means,

wherein said throttling means intercepts a packet on a TCP/IP connection between the first apparatus and the second apparatus, and

wherein said throttling means effects the throttling by modifying the packet to cause the second apparatus to change an amount of data it sends before awaiting a TCP ACK from the first apparatus.

FCHS_WS 2181007v1

EXHIBIT 2

DirecPC™

HUGHES
NETWORK SYSTEMS
A HUGHES ELECTRONICS COMPANY

DirecPC
Turbo-Throttling
Feature Design Document

Document HNS-11422

Version 1.0



**Hughes Network Systems, Inc.
11717 Exploration Lane
Germantown, Maryland 20876**

(301) 428-5500

© Copyright Hughes Network Systems, Inc. 1996.

All rights reserved. This Publication and its contents are proprietary to Hughes Network Systems, Inc. No part of this publication may be reproduced in any form or by any means without the written permission of Hughes Network Systems, Inc., 11717 Exploration Lane, Germantown, Maryland 20876.

Revision History

Revision	Date of Issue	Author	Scope
0.1	[REDACTED]	Vivek Gupta	Initial draft
0.2	[REDACTED]	Vivek Gupta	Incorporated Doug's comments
1.0	[REDACTED]	Vivek Gupta	Modify for new window-sizing based Turbo-Throttling

Table of Contents

1. Introduction	7
1.1 Purpose	7
1.1.1 Intended Audience	7
1.2 Scope.....	7
1.3 Definitions and Acronyms	7
1.4 References	7
1.5 Font Selection.....	7
1.6 Open Issues.....	7
1.7 Requirements Traceability Matrix	7
2. System Overview.....	9
2.1 Basic Throttling concept	9
2.2 Running-Average Throughput (RAT).....	9
2.3 New data structures at the Hybrid Gateway	9
2.3.1 Per-service plan parameters.....	9
2.3.2 Per-user statistics	10
2.4 User Modes	11
2.4.1 User Mode transitions	11
2.4.2 Hysteresis	14
2.5 Maintaining connect minutes	14
2.5.1 Connect Messages	14
2.5.1.1 Updates to DAK 1.6.....	14
2.5.1.2 Updates to HGW	15
2.6 New Hybrid Gateway Statistics.....	15
2.7 ACS/Oracle Database enhancements.....	15
2.7.1 Updating Service Plan information.....	15
2.7.2 Changes to HGW reconfiguration process	15
2.7.3 New triggers	16
2.7.4 Updates to Oracle database tables.....	16
2.7.5 New forms.....	16
2.7.6 New reports.....	16

List of Tables

Table 1. Definitions and Acronyms.....	7
Table 2. Requirements Traceability Matrix.....	7

Chapter 1

Introduction

1.1 Purpose

This document details the design for the DirecPC Turbo Internet Statistics feature.

1.1.1 Intended Audience

This document is intended for use by software engineering and test personnel.

1.2 Scope

This document details the design of all software affiliated with DirecPC Turbo Internet Statistics feature.

1.3 Definitions and Acronyms

This section provides definitions for all acronyms and terms introduced in or unique to this document.

Table 1. Definitions and Acronyms

Term	Description
Customer	The end user of the DirecPC software. The person initiating the auto commissioning process
NOC	DirecPC Network Operations Center
DirecPC Access Kit (DAK)	The hardware and software that constitutes the Remote DirecPC product

1.4 References

[1] DirecPC Turbo-Throttling Feature Requirements Specification (HNS-11232)

1.5 Font Selection

The first occurrences of all **variables/data structures** are marked with a Bold-Italic font, the later occurrences of the **variables** are marked with Italic font. The values (which may be taken by variables) are underlined.

1.6 Open Issues

1.7 Requirements Traceability Matrix

The following table is a synopsis of requirements from the DirecPC Turbo-Throttling requirements with a cross reference to which design-item verifies compliance to the requirement.

Table 2. Requirements Traceability Matrix

Requirement ID	Section ID
[R:0019]	2.2
[R:0033]	2.3.1
[R:0021]	
[D:0022]	

Requirement ID	Section ID
[R:0019]	2.2
[R:0033]	2.3.1
[D:0023]	2.5
[R:0024]	2.6
[R:0025], [D:0026]	2.3.2
[R:0027]	Error! Reference source not found.
[R:0029]	Error! Reference source not found.
[R:0030]	Error! Reference source not found., 2.3.1
[R:0031]	Error! Reference source not found.
[R:0012]	2.7.1
[R:0013]	2.7.1, 2.7.2, 2.7.3
[R:0014]	2.7.5

Chapter 2 Design

2.1 Basic Throttling concept

Turbo-throttling is based on the fact that the user will be assigned an initial window size and a bucket capacity based on the amount of average data rate that is allowed by the service plan to which he belongs. If the user goes beyond the average data rate (computed by counting the number of bytes transferred by the user in the last hour, for example), he will be "throttled" through a reduction in his window size and some non-DNS UDP discards.

Two levels of throttling will be maintained and each level will have its corresponding window size assignment and UDP discard rate.

Hysteresis will be introduced while moving a user from the throttled state to a non-throttled one, so that he does not oscillate between the two states.

Finally, the feature also includes a means to disable data forwarding (or download) to a user via the NOC if the user is not logged on (that is, if no data has been received from the user in the past minute)

2.2 Running-Average Throughput (RAT)

The average data rate or the running average throughput will be maintained using the leaky-bucket approach used for rate-based flow control in Frame Relay.

The maximum burst size that can be tolerated at any given time depends on the bucket size "AND" on the current available capacity of the bucket. Hence, this scheme implicitly works well for an internet surfer who downloads bulk data (at a fast rate) and then spends some time reading the data or searching for other stuff (while staying connected!). If the user attempts to download big chunks of data and disconnect and again download big chunks, this scheme will put him in the throttle modes

The RAT will be stored as part of Enhanced TI statistics and will also be accessible through the *hwstat* utility.

2.3 New data structures at the Hybrid Gateway

The following data structures will be maintained at the Hybrid Gateway (some of them will be configured using the ACS). These data structures have been defined here in order to aid the explanation of the design concepts that follow.

2.3.1 Per-service plan parameters

The Hybrid Gateway will maintain the following new (configuration) parameters on a per-service plan basis:

- **Throughput Soft Threshold** (in kbps)
- **Throughput Hard Threshold** (in kbps)
- **Running Average Duration** (in minutes)
- **WSizeMax** (in bytes)
- **WSizeSoftThresh** (in bytes)
- **WSizeHardThresh** (in bytes)
- **UDP Discard Soft Threshold** (frequency of packet discard)
- **UDP Discard Hard Threshold** (frequency of packet discard)

- **BucketDepthSoftThresh** (= $\{ \text{Throughput Soft Threshold} * \text{Running Average Duration} \}$ bytes). If the user's bucket depth is less than this value, he will not be throttled, that is, his window size will be WMax and none of his UDP traffic will be discarded. If the user's bucket depth exceeds this value, he will enter a mode of Soft-Throttling wherein his window size will be reduced to WSizeSoftThresh and his (non-DNS) UDP traffic will be discarded at the rate of UDPDiscardSoftThresh
- **BucketDepthHardThresh** (= $\{ \text{Throughput Hard Threshold} * \text{Running Average Duration} \}$ bytes). If the user's bucket depth exceeds this value, he will enter a mode of Hard-Throttling wherein his window size will be reduced to WSizeHardThresh and his (non-DNS) UDP traffic will be discarded at the rate of UDPDiscardHardThresh
- **Leak Per Minute** (= $\{ \text{Throughput Hard Threshold} * \text{seconds per minute} \}$ bytes): Number of bytes which will be "leaked" (or subtracted) from the user's bucket at the minute boundary. Since the **Leak Per Minute** is based on the **Throughput Hard Threshold**, it allows a faster leak than the data rate which the user may be able to sustain in the Soft-Throttle state (so as to allow the user to return to a non-Throttled state).
- **Excluded Daily Minutes**: Number of minutes every day (after 00:00:00 am, EST) that a user will be allowed to operate with Maximum window size (WMax) and no UDP discards. In this duration, the user's traffic will not be included for RAT calculation.
- **Hysteresis Duration** (in minutes): This decides the amount of time that a user will be forced to stay in the Soft-Throttled mode after his RAT has fallen below the **Throughput Soft Threshold**. No Hysteresis is planned to be used in the transition from the Hard-Throttled state to Soft-Throttle
- **Enable Disconnect-Forward**: If set, this will allow forwarding of data to a user who is not connected.

The following points may allow a better understanding of the parameters:

1. Running Average Duration affects both the bucket depth thresholds
2. Throughput Hard Threshold affects both *bucket depth hard threshold* and bucket leak rate
3. Throughput Soft Threshold affects only the *bucket depth soft threshold*

2.3.2 Per-user statistics

The Hybrid Gateway will maintain some new (non-volatile) statistics on a per-IP address basis: (all these statistics will be stored as part of the billing data structure already maintained in the HGW)

1. Non-persistent statistics: (these need not be preserved across HGW restarts)
 - Hourly statistics:
 - **active**: Boolean value indicating whether the user was active or not in the last minute
 - **NumConnMsgs**: Number of Connect Messages (explained in a later section) received in the past hour from the remote PC.
 - **Hourly connect minutes**: number of minutes in which some data was received from or transmitted to the remote PC.
 - **Hourly Throttled minutes**: amount of time the user was in one of the throttling states in the past hour
 - **NumUDPDiscards**: Number of UDP packets discarded because of throttling in the past hour
 - **IHpktDiscards**: Number of packets (received from IHs) discarded because of user being disconnected
 - Other statistics:
 - **TIServiceId** (service plan to which the user belongs)

- **ConnectMode** (valid states are mentioned in section 2.3): Current mode/state of user.
Initialized to DISCONN
- **NegPPPIPaddr**: PPP-IP address negotiated by the remote PC with the ISP
- **Wsize**: Current window size being assigned to all TCP connections opened by remote PC. It should be understood that this window size will "always" be less than the FCTWS (which, in turn, is less than the MCTWS).

2. Persistent statistics:

- **Daily connect minutes**: used to provide a better initial experience (provided $X > 0$). The "start" and "end" of a day will be defined by the HGW's local clock (regardless of the location of a particular user).
- **Current Bucket size** (number of bytes in bucket): This decides the user mode (throttled/non-throttled)
- **LastMode**: this is needed to be able to start in the last mode the user was in before he disconnected or before the HGW was restarted. Initialized to XSAT (so that the first time the user uses DirecPC, he starts off in the XSAT mode)
- **HysteresisMin**: This counts the number of minutes the user has been in the Recovery mode (explained later)
- **SuccUDPtxed**: This statistic is used only in the throttled states and counts the number of successive UDP packets that have been transmitted (that is, without being discarded)

2.4 User Modes

Each user can be in one of the following modes at any instant of time:

1. DISC (Disconnected): the user is disconnected. This issue is very tricky and the following statement may clarify the matter. If no data transmission takes place to/from the remote in a particular minute, the user state is set to DISCONN. From then on,
 - if *EnableDisconnectFwd* is NOT set for the user, any "new data" received from the internet will NOT be forwarded unless some data is received from the remote OR some re-transmission of TCP data (which may have been queued up for the remote PC) takes place. In either of these cases, the user mode will be changed to the "LastEffectiveMode"
 - if *EnableDisconnectFwd* is set for the user, data transmission will take place leading to change in mode to "LastEffectiveMode".
1. XSAT (Exclusive-Satellite): No-Throttling (Params: WMax, 0), User's traffic is NOT used in running average throughput (RAT) calculation.
2. ISAT (Inclusive-Satellite): No-Throttling (Params: WMax, 0), User's traffic is included for RAT calculation
3. STHR (Soft-Throttle): 1st level of Throttling (Params: W1, U1), User's traffic is included for RAT calculation
4. RECO (Soft-Throttle-Recover): Recovery from STHR (Params: W1, U1), User's traffic is included for RAT calculation
5. HTHR (Hard-Throttle): 2nd level of Throttling (Params: W2, U2), User's traffic is included for RAT calculation.
6. INVL (Invalid): Software error (should never happen)

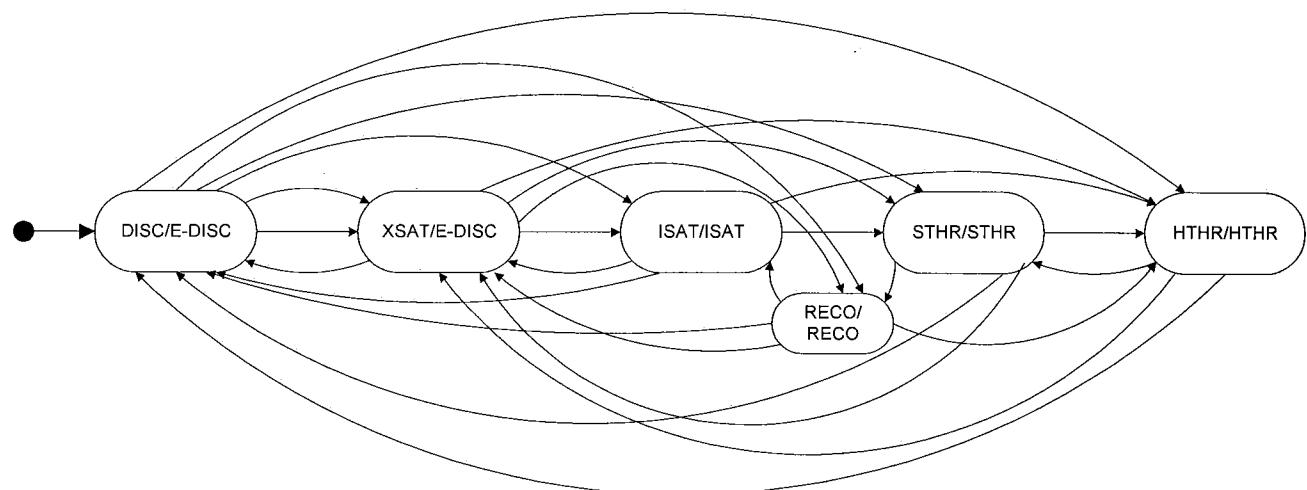
2.4.1 User Mode transitions

The following criteria decide the transition from one user mode to another:

1. In a simplistic view of transitions, a user moves from DISC to XSAT when he first logs on. In the XSAT mode, the user is assigned a window size of WMax and his data transfers are not included in RAT calculation. After staying in the XSAT mode for DailyExcludedMinutes, the user moves to ISAT mode wherein the user's window size remains the same; however, his data transfers are used for RAT calculation.
2. All user mode transitions except those from DISC to other modes, are done at the minute boundary
3. DISC to any other mode transition is done as soon as data is received/sent to the user
4. No window size adjustments need to be made when making a transition from DISC to any mode (other than XSAT). When moving from DISC to XSAT, the window should be set to WMax.
5. When moving from XSAT to any mode, window size needs to be adjusted based on LastMode! In particular, window size adjustment is needed if moving to STHR, RECO and HTHR.

Each state depicts:
CurrentMode/LastMode

E- Any state except-
ANY Any state



2.4.2 Hysteresis

Hysteresis is provided by two means:

- a) The user is allowed to go beyond his allotted quota by an amount (whose maximum value is equal to the number of bytes that can be downloaded in a minute).
- b) The user is not switched back to the ISAT mode unless he has stayed for *HysteresisMin* minutes in STHR mode "after" his RAT has fallen below *Throughput Soft Threshold*.

2.5 Maintaining connect minutes

In the present Turbo-Internet model, the Billing records only include "active" minutes, that is, the minutes in which some data transfer was done through the Hybrid Gateway on behalf of the particular remote host. This scheme will need to be changed with the "Hourly" service plans which require the knowledge of "connected" minutes rather than "active" minutes at the Hybrid Gateway.

The connect minutes maintenance requires a periodic send of messages (hereinafter, referred to as "Connect Messages") from the DAK to the HGW to update the "connect minutes" at the HGW. It also requires a periodic send of "status messages" from the HGW to all connected remotes so that the user can be notified about the number of minutes/hours used in the current month.

2.5.1 Connect Messages

2.5.1.1 Updates to DAK 1.6

As long as the modem is connected, Turbo-Internet will send out periodic messages to the Hybrid Gateway. In order to comply with Internet standards, the protocol chosen for this communication is HMP (Host-Monitoring Protocol) [RFC 869]. These messages will be sent in a tunneled mode (since the IP address of the remote will not be available in the un-tunneled mode) to the Hybrid Gateway.

The default destination for the tunneled messages will be the DNS Server (explained in later section) and the default time-period for message send will be 20 seconds. The default will be changeable through a registry key (maximum value=60):

HKEY_LOCAL_MACHINE\SOFTWARE\Hughes Network Systems\DirEcPc\TurboInternet\
TimerConnMsgSend

An adhoc addition of periodic message send to Turbo-Internet completely breaks the Idle-timeout handling in TAPIDLO (which is supposed to time-out if the user is not active for a configurable amount of time). In order to get around this problem, all messages which are sent out to Internet via TurboInternet have been classified as Internal and External messages. Internal messages are generated from within Turbo-Internet and should not be accounted as part of user activity.

Since there is no mechanism to notify TAPIDLO about the message type, there are two alternatives:

- a) modify the current TAPIDLO functions (specifically, TapiDloSend) to accept an extra parameter
- b) add a new TapiDlo function.

Since the former choice necessitates modifying all DLLs which use TapiDloSend, the latter choice will be adopted. This will ensure backward compatibility. The new function will be called TapiDloSendEx and will accept an extra parameter (unsigned int) which will indicate the message type. The queues and data structures used in PPP and TurboInternet software will need to be modified to differentiate between the two message types.

The connect messages will be defined on top of HMP (Host Monitoring Protocol) and will use the following fields for each data item: type (1 byte), length (1 byte), data (length bytes). The following types will be defined initially:

type = 0x1 (connect-speed), length = 0x2, data = 0x2580 (9600 bps)

On receiving this message, the user's *Terrestrial Bandwidth* should be set to the value as indicated in the message.

It should be noted that the connect messages will be sent "only if" the user is in Satellite-only or Selective Terrestrial. No messages will be sent in the Terrestrial-only mode.

2.5.1.2 Updates to HGW

On receipt of a connect message, the HGW should verify the checksum, update HMP statistics for the HGW as a whole and HMP statistics for the user and also update the active status of the user (this will result in updation of active-minutes at the end of the minute). The HGW should also update the connect speed for the remote. The connect message should, thereafter, be dumped at the HGW.

The Release 1.6 and earlier HGWs will treat these messages as "Other IP" messages and will forward them to the "tunneled destination" after stripping off the outer IP header. If the tunneled destination is set to the HGW itself, the message will appear once more on the HGW LAN segment (which is not desirable). Hence the tunneled destination has been set to the DNS Server which, in turn, will discard the unsolicited message. With a Release 1.6 or later remote, the active-minutes reported will be equal to the "connected minutes".

2.6 New Hybrid Gateway Statistics

The user-specific statistics accessible through the hgwstat utility will be enhanced to report the current user mode, the bytes that have been sent terrestrially during the last hour and the monthly connect minutes. The user-statistics will also carry information about the number of Connect messages received from the remote PC in the last hour.

The hourly statistics will be enhanced to include the number of bytes sent over the terrestrial and satellite path, and the number of minutes spent in either mode. The per-application statistics will, however, not be classified under terrestrial/satellite modes.

Since the monthly connect minutes for each user have to be maintained across HGW restarts, they will be dumped to a file on an hourly basis and on every software termination. This file will be named usage.sta and will have records in the form SSSSSSS=BBBBBBBBBB;XX;YY;ZZ;AA (where SSSSSSS is the site-id, BBBB BBBB BBBB is the *current bucket size*, XX denotes the *LastMode*, YY denotes the *DailyConnMin*, ZZ denotes the *HysteresisMin*, AA denotes the *SuccUDPtXed*). This file will reside in the current "run" directory.

At startup, the HGW will look for this file. If the file is not found, it will post an error message in the log file. Else it will initialize the *monthly connect minutes* and the *current bucket size* for each user using this file.

2.7 ACS/Oracle Database enhancements

2.7.1 Updating Service Plan information

A utility will be provided which will take a file (with complete path) as a parameter and will process the user-profile provided in an ascii file (each line of which contains a 8 character site-id and a 2-digit numeric service plan-id, both separated by a single space) and update the ACS database accordingly. It will be the responsibility of the Network Operations Center to work with the Billing subsystem to obtain this file periodically (say, through a cron job which uses FTP to get the file). The format of the file will be added to the DirecPC Product Technical Specification document (Pubs No. 8050134).

2.7.2 Changes to HGW reconfiguration process

The acshconf program will be changed to add the following items to the hybridgw.cfg file:

- TIServicen = T1;T2;D;WMax;W1;W2;U1;U2;H;X;F (where T1, T2, D, WMax, W1, W2, U1, U2, H, X and F are parameters as defined in [1]), where n would take values from 1 to 24 (for 24 possible Service Plans for now).

The acshconf program will also be changed to add the ServicePlanId field to each line in the individual range files.

2.7.3 New triggers

There should be triggers based upon the following actions:

- Changes to the TurboInternetServices table (that is, change in any service-plan parameters)
- Changes to the TIServiceld field of remotes table

2.7.4 Updates to Oracle database tables

The following Oracle tables need to be modified/added:

- Hybridgws: This will have the following additional field: R (Minimum Terrestrial redirect). The "UseTokenRing" field will now be used to decide between a FastEthernet or TokenRing interface to the Satellite Gateway and CAC.
- TurboNetServices: This will have the new fields as defined in [1] that is, T1, T2, D, WMax, W1, W2, U1, U2, X, F
- TIStats: This will have the following additional fields (each record is an hourly record):
 - Number of Hourly Connect minutes
 - Current Bucket size

2.7.5 New forms

A form will be added to allow updatations to the Service Plan parameters. These parameters are defined as part of TurboNetServices Table and will be changed by the NOC operator in order to change a service plan definition.

2.7.6 New reports

The reports based on Turbo-Internet statistics will be enhanced to include Terrestrial bytes and minutes and to include Average Bucket size for each user.

EXHIBIT 3

DirecPC™ Turbo-Throttling

Data Sheet

The DirecPC Turbo-Throttling feature provides the additional capabilities to DirecPC Turbo-Internet to support an enforced running average fair use policy and also performance-differentiated levels of service. The highlights of Turbo-Throttling include:

- Maintaining a running average of the user's total throughput. The running average calculation is based on the leaky-bucket approach used for rate-based flow control in ATM and Frame Relay. The running average is maintained across user sessions.
- Assigning an initial per-connection TCP window size WMax to each user and clamping it to W1 or W2 when the running average throughput exceeds thresholds T1 and T2 ($T1 \leq T2$). WMax, W1 and W2 will each be greater than the minimum segment size (536 bytes) and $WMax \geq W1 \geq W2$.
- Optionally, providing two levels of (non-DNS) UDP discard, U1 and U2, linked with T1 and T2 respectively ($1 \leq U2 \leq U1 \leq 1000$). If T1 is reached, every $U1^{th}$ (say, 50th) UDP packet will be dropped and if T2 is reached, every $U2^{th}$ (say, 1st that is, every) packet will be dropped.
- Optionally giving users a better initial experience, by having the HGW not include the first X minutes (default = 0) of activity each day in the running average. During this time the user may run at whatever rate he can achieve.
- Provide a means to enable/disable data forwarding when a user is not connected (F: enable disconnect forwarding) on a service plan basis.
- Modifying DAK and HGW software so that active minutes in billing records equal connected minutes.
- Updated auto-commissioning and Turbo-Internet statistics
- Modifying Auto-commissioning software to allow the configuration of T1, T2, D, WMax, W1, W2, U1, U2, F and X on a service plan basis. The HGW reconfiguration process will also be modified to include the service-plan information for each user.

The Per-Service Plan Turbo-Throttling parameters are as follows:

- D: Duration over which Running Average Throughput is calculated, default 60 minutes.
- WMax: Maximum window size for a TCP connection - used to limit the service plan's relative performance (single connection throughput), 48 kbytes is roughly equivalent to 400 kbps, 24 kbytes is roughly equivalent to 200 kbps.
- T1, T2: Throughput Thresholds - increasing throughput clamping is employed as the running average throughput exceeds these values. Defaults, T1 = 40 kbps, T2 = 64 kbps.
- W1, W2: Clamped window sizes when running average throughput exceeds T1 and T2 respectively. For example E1 = 2000 bytes and W2 = 600 bytes.
- U1, U2: Defines the UDP discard rate when the running average throughput exceeds T1 and T2 respectively. For example, U1 = 50, that is, 1 out of 50 and U2 = 1, that is, 1 out of 1 (every packet).
- F, Enable Disconnect-Data forwarding
- X, Daily Initial Un-averaged Minutes - non-zero ensures a good initial experience for the user.

EXHIBIT 4

Harvey Lindenbaum
02/09/98 04:36 PM

To: John Kenyon/HNS@HNS, Barbara Stavely/HNS@HNS, Kelly Petronic/HNS@HNS, William Chao/HNS@HNS, Doug Dillon/HNS@HNS, iFaenson, Jerry Longo/HNS@HNS,
cc: Ramesh Belani/HNS@HNS, Satyajit Roy/HNS@HNS, murr, Vaishali Sesha/HNS@HNS, Roderick Ragland/HNS@HNS, Tsanchi Li/HNS@HNS, Gabriel Olariu/HNS@HNS, Anil Vohra/HNS@HNS, msyed, Trung Tran/HNS@HNS, Matthew Baer/HNS@HNS, Rajeev Kubba/HNS@HNS

Subject: DirecPC Weekly Report

ENGINEERING WEEKLY REPORT BY PRODUCT LINE

WEEK ENDING: 2/6/98

PROJECT LEADER: Harvey Lindenbaum

PRODUCT LINE: DirecPC JOB NUMBER: 12542, 12543

PRODUCT: DirecPC

RED FLAG ITEMS: None

MAJOR ACCOMPLISHMENTS/MILESTONES:

Turbo Throttling 2
FIT testing of Turbo Throttling 2 has been completed. It has been turned
over to
Ramesh Belani for Release Testing.

EXHIBIT 5

DirecPC™



DirecPC
Turbo-Throttling
Feature Requirements Specification

Document HNS-11232

Version 1.1

Hughes Network Systems, Inc.
11717 Exploration Lane
Germantown, Maryland 20876

(301) 428-5500



© Copyright Hughes Network Systems, Inc.

All rights reserved. This Publication and its contents are proprietary to Hughes Network Systems, Inc. No part of this publication may be reproduced in any form or by any means without the written permission of Hughes Network Systems, Inc., 11717 Exploration Lane, Germantown, Maryland 20876.

- ® The following are registered trademarks:
- DirecPC

Revision History

Revision	Date of Issue	Author	Scope
0.1	[REDACTED]	Vivek Gupta	Initial draft
0.2	[REDACTED]	Vivek Gupta	Incorporated inspection comments
0.3	[REDACTED]	Vivek Gupta	Redo draft based on new (Hourly) Service Plan definition
0.4	[REDACTED]	Vivek Gupta	Added a requirement and incorporated changes suggested in Inspection
0.5	[REDACTED]	Vivek Gupta	Incorporated changed data sheet, incorporated Doug's comments (Highest ReqId: 0032)
0.6	[REDACTED]	Vivek Gupta	Incorporate new data sheet, drop terrestrial redirect and use window sizing for differentiated levels of service (Highest ReqId: 0035)
1.0	[REDACTED]	Vivek Gupta	Changes for Turbo-Throttling II (added data sheet, modified Reqs 0033, 0034; added Reqs 0036-0043)
1.1	[REDACTED]	Vivek Gupta	Added data sheet for configuring Turbo-Throttling II parameters

Table of Contents

1. Introduction.....	1
1.1 Purpose.....	9
1.1.1 Intended Audience.....	9
1.2 Scope	9
1.3 Terminology.....	9
1.4 Definitions and Acronyms	10
1.5 References	11
1.6 Open Issues	11
2. Overview	16
2.1 Overview	16
2.2 Deliverables.....	16
2.3 Compatibility and Expected System Impact.....	16
3. Functional Requirements.....	19
3.1 Functional Requirements	19
3.1.1 [R:0019] Running Average Throughput.....	19
3.1.2 [R:0033] Assign initial window size based on Service Plan.....	19
3.1.3 [D:0022] Exclude first few minutes of usage every day from running average calculation	19
3.1.4 [R:0034] Clamp window size based on running average throughput	19
3.1.5 [R:0036] Per-user Throughput Throttling.....	19
3.1.6 [R:0035] Provide two levels of UDP discard corresponding to T1 and T2	20
3.1.7 [R:0037] TCP Connection Limiting	20
3.1.8 [R:0038] Bucket Leak	20
3.1.9 [D:0023] Use “connected” minutes for active minutes and running average calculation.....	20
3.1.10 Persistent data at HGW.....	20
3.1.11 [R:0031] Disable data forwarding to a user when not logged on.....	20
3.1.12 [R:0012] Update of Service Plan information for each user	20
3.1.13 [R:0039] Unlimited number of service plans	21
3.1.14 [R:0013] HGW reconfiguration	21
3.1.15 [R:0042] Enhanced TISstats	21
3.2 Non-Throttling requirements	21
3.2.1 [R:0040] Local Key Table	21
3.2.2 [R:0041] Change in Flow Control (FC) mechanism.....	21

3.3 Obsolete Requirements	21
3.3.1 [X:0020] Terrestrial redirection of user traffic	21
3.3.2 [X:0021] Return to satellite path	22
3.3.3 [X:0024] Updated Turbo-Internet Statistics	22
3.3.4 [X:0027] Terrestrially redirect some portion of traffic at all times.....	22
3.3.5 [X:0028] Minimum session time.....	22
3.3.6 [X:0029] Hourly usage notification to user.....	22
3.3.7 [X:0030] Per-Service Plan, Per-HGW Turbo-Throttling parameters.....	22
3.3.8 New configuration parameters for Hybrid Gateway	23
3.3.9 New statistics at Hybrid Gateway and in Oracle Database	23
3.3.10 [X:0010] New priority levels at HGW and SGW.....	23
3.3.11 [X:0011] Terrestrial return of Turbo-Internet traffic.....	24
3.3.12 Hour-meter display	24
3.4 [R:0032] User Interface	24
3.4.1 [R:0014] Configuration of Turbo-Throttling parameters via Oracle Forms	24
3.5 External Interface Requirements.....	24
3.5.1 User Interfaces/Characteristics	24
3.5.2 Hardware Interfaces	24
3.5.3 Software Interfaces.....	24
3.5.4 Communications Interfaces	24
3.5.5 Network Management Interfaces.....	24
3.6 Expected Enhancements	24
4. System Attributes	25
4.1 Performance.....	25
4.1.1 [R:0015] No impact on existing Billing/Performance statistics.....	25
4.2 Capacity and Resources	25
4.3 Scalability	25
4.4 Design Constraints	25
4.4.1 [R:0016] Standards Compliance.....	25
4.4.2 Hardware Limitations	25
4.4.3 Physical Environment Considerations	25
4.4.4 Timing Constraints.....	25
4.4.5 [R:0017] Software Compatibility/Environment Considerations	25
4.4.5.1 [R:0043] Backwards Compatibility	25
4.4.6 Unit Cost.....	26

4.5 Time to Market	26
4.6 Testability	26
4.7 Attributes	26
4.7.1 Availability	26
4.7.1.1 Startup/Recover Time	26
4.7.1.2 Fault Tolerance	26
4.7.1.3 Mean Time Between Failure (MTBF)	26
4.7.2 Security	26
4.7.3 Privacy	26
4.7.4 Debugging/Trace Capability	26
4.7.5 [R.0018] Maintainability	26
4.7.6 Adaptability	27
4.7.7 Portability	27
4.8 Installation Requirements	27
4.9 Customization Requirements	27
4.10 Other Requirements	27

List of Tables

Table 1. Requirements Criteria.....	10
Table 2. Definitions and Acronyms	11

Chapter 1 Introduction

1.1 Purpose

The purpose of this document is to define the requirements for the DirecPC Turbo-Throttling feature. This document has several goals:

- Establish a basis for agreement between Engineering and various operating divisions of HNS about what the product must do to meet the needs of the customer and those who must manufacture, support, and maintain it;
- Provide a basis for generating a project plan, estimating costs and schedules;
- Provide a basis for validation and verification of the final product;
- Provide a basis for maintenance and enhancement of the product.

1.1.1 Intended Audience

This document is applicable to all of the HNS operating divisions that will fund, implement, manufacture, sell, use, maintain, and support this capability such as:

- Hardware and Software Engineering
- Production, Quality Assurance, and System Test
- Field Services (Installation and Maintenance)
- Customer Services - Training, Support, Documentation
- Program Management
- Marketing
- Applications Programmers
- Legal

At the discretion of Program Management, some portions of this document may be disclosed to prospective customers.

1.2 Scope

This document defines the requirements for the Turbo-Throttling feature as it affects the DirecPC NOC and Remote software.

1.3 Terminology

It is expected that all requirements and capabilities listed in this document will be implemented for the current release and all future releases of this system. Each requirement in this document must be identified using one of the notations described in Table 1:

Table 1. Requirements Criteria

[R:nnnn]	Indicates meeting this requirement is required
[O:nnnn]	Indicates meeting this requirement is an objective . An objective is within the current scope of the features, but may be sacrificed to meet schedule
[D:nnnn]	Indicates meeting this requirement is desirable . A desirable requirement is one which would be nice to have, but has been prioritized as not being important enough to actually build. A desirable requirement is not within scope (will not be built).
[X:nnnn]	Indicates a requirement which was suggested at one time, but which has already been decided as not something we are at all interested in doing. X:nnnn requirements are sometimes left in a document to document that a requirement has been specifically rejected

Together, the letters **R**, **O**, and **D**, and the **nnnn** portion of the notation form a Document-Unique Requirement ID (DURID) that uniquely identifies the requirement. The **nnnn** portion of the DURID can be any numeric value as long as it, together with the letter **R**, **O**, or **D** form an ID that is unique across the whole document. The **document number + DURID** form a primary key used to identify a specific requirement in a requirements traceability matrix.

The following is a list of constraints related to the use of DURIDs:

- DURIDs must be enclosed in square brackets [] and only one DURID must be enclosed in a set of brackets. This is to ensure consistent markup and facilitate future automated requirement lookup schemes.
- DURIDs should be placed after paragraph numbers and at the beginning of each requirement listed.
- Once a DURID is assigned to a requirement, it must never change.
- DURIDs identify a single requirement for the life of that requirement.
- DURIDs should not be reused when a requirement is deleted if a requirements traceability matrix already exists for the feature described by this document.
- DURIDs can be duplicated if they are in different documents.
- The **nnnn** portion of the DURID should be unique across the document without regard to the DURID letter portion. This allows a document writer to more easily determine the next available DURID number because the combination of letter + number does not need to be taken into account. This would also enable development of an automated DURID numbering macro, for instance, to be simplified.

The following are examples of requirements identified using the DURID notation:

- 3.1.2. [R:0001] A Requirement To Do This Thing
- 3.1.2. [O:0002] This Thing Should Be An Objective
- 3.1.3. [D:0003] A Desire To Do This Other Thing

1.4 Definitions and Acronyms

This section provides definitions for all acronyms and terms introduced in or unique to this document.

Table 2. Definitions and Acronyms

Term	Description
Customer/User	The end user of the DirecPC software
NOC	DirecPC Network Operations Center
DirecPC Access Kit (DAK)	The hardware and software that comprises the remote DirecPC product
HGW	Hybrid Gateway

1.5 References

- [1] DirecPC Turbo-Internet Statistics Feature Requirements Specification
- [2] DirecPC Turbo-Internet Statistics Feature Design Specification

1.6 Open Issues

- There are no open issues

DirecPC™ Turbo-Throttling

Data Sheet

The DirecPC Turbo-Throttling feature provides the additional capabilities to DirecPC Turbo-Internet to support an enforced running average fair use policy and also performance-differentiated levels of service. The highlights of Turbo-Throttling include:

- Maintaining a running average of the user's total throughput. The running average calculation is based on the leaky-bucket approach used for rate-based flow control in ATM and Frame Relay. The running average is maintained across user sessions.
- Assigning an initial per-connection TCP window size WMax to each user and clamping it to W1 or W2 when the running average throughput exceeds thresholds T1 and T2 ($T1 \leq T2$). WMax, W1 and W2 will each be greater than the minimum segment size (536 bytes) and $WMax \geq W1 \geq W2$.
- Optionally, providing two levels of (non-DNS) UDP discard, U1 and U2, linked with T1 and T2 respectively ($1 \leq U2 \leq U1 \leq 1000$). If T1 is reached, every $U1^{\text{th}}$ (say, 50th) UDP packet will be dropped and if T2 is reached, every $U2^{\text{th}}$ (say, 1st that is, every) packet will be dropped.
- Optionally giving users a better initial experience, by having the HGW not include the first X minutes (default = 0) of activity each day in the running average. During this time the user may run at whatever rate he can achieve.
- Provide a means to enable/disable data forwarding when a user is not connected (F: enable disconnect forwarding) on a service plan basis.
- Modifying DAK and HGW software so that active minutes in billing records equal connected minutes.
- Updated auto-commissioning and Turbo-Internet statistics
- Modifying Auto-commissioning software to allow the configuration of T1, T2, D, WMax, W1, W2, U1, U2, F and X on a service plan basis. The HGW reconfiguration process will also be modified to include the service-plan information for each user.

The Per-Service Plan Turbo-Throttling parameters are as follows:

- D: Duration over which Running Average Throughput is calculated, default 60 minutes.
- WMax: Maximum window size for a TCP connection - used to limit the service plan's relative performance (single connection throughput), 48 kbytes is roughly equivalent to 400 kbps, 24 kbytes is roughly equivalent to 200 kbps.
- T1, T2: Throughput Thresholds - increasing throughput clamping is employed as the running average throughput exceeds these values. Defaults, T1 = 40 kbps, T2 = 64 kbps.
- W1, W2: Clamped window sizes when running average throughput exceeds T1 and T2 respectively. For example E1 = 2000 bytes and W2 = 600 bytes.
- U1, U2: Defines the UDP discard rate when the running average throughput exceeds T1 and T2 respectively. For example, U1 = 50, that is, 1 out of 50 and U2 = 1, that is, 1 out of 1 (every packet).
- F, Enable Disconnect-Data forwarding
- X, Daily Initial Un-averaged Minutes - non-zero ensures a good initial experience for the user.

DirecPC™ Turbo-Throttling II

DataSheet

DirecPC Turbo-Throttling II is intended to provide a stricter enforcement of the Turbo-Internet fair use policy and to prepare for ISP Bundling. The highlights of Turbo-Throttling II include:

- TCP connection limiting: Defines the number of simultaneous TCP connections that can be opened by a user.
- Per-user throughput throttling: When throttled, the user's bandwidth will be shared among the TCP connections. This will ensure that, even with multiple connections, a user's aggregate throughput will not exceed the configured value. The bandwidth will be shared by dividing the window size equally among all "bulk-transfer" type connections.

(Note: A connection is classified as a "bulk-transfer" connection if, in the *recent past*, the number of unacknowledged bytes exceeded a configured value, typically 1kbytes. Otherwise, the connection is classified as an "interactive" connection. It may be noted that all connections start off as "interactive" connections and may become "bulk-transfer" later).
- Change in Flow Control (FC) Mechanism: Under high load conditions, the flow control mechanism throttles the window size of each TCP connection to the *same value*. This has led to a condition wherein a hard-throttled user's TCP window is not affected at all, while a non-throttled user's TCP window may shrink to a size equal to the hard-throttle window size. The new FC mechanism will maintain a FC meter (with values from 1% to 100%); the meter value will be computed based on the load conditions on SGW. A user's effective window size will be determined by multiplying this percentage with the window size derived on the basis of his throttled state.
- Bucket leak: New parameters will be added to specify the bucket leak when connected and when disconnected. The latter leak rate will be employed especially for users with ISP Bundling service plans, so that they do not stay connected at HNS' expense just to reduce their running average throughput.
- New Per Service-Plan parameters (added to TurboInetServices table):
 - Max TCP Connections, Connected Leak rate (kbps), Disconnected Leak rate (kbps)
 - Peak-UnThrottled Throughput (kbps), Peak-SoftThrottled Throughput (kbps), Peak-HardThrottle Throughput (kbps): These, in turn, will yield window sizes (with RTT being set on a per-HGW basis)
- Enhanced TIStats: The following fields will be added (on a per-user basis) to the TIStats table:
 - Number of TCP connections rejected, Number of Hard Throttled Minutes, MaxBulkTransferConns
 - Service Plan of the user, Hybrid Gateway Id, Bytes Received

The following HGW enhancements will also be bundled with Turbo-Throttling II:

- Local Key Table: The HGW maintains a table to store the encryption keys for each remote. This table will now store key information only for the remotes configured on the particular HGW.
- Unlimited number of service plans: The last release of HGW (HGNT1.8.4) limited the service plans to 30. The new release (HGNT2.x) will not restrict the number of service plans.

Backward compatibility:

- Oracle database upgrade with no HGW upgrade: HGW will ignore the new parameters
- HGW upgrade with no database upgrade: HGW assumes version 1 of Throttling with MaxTCPConns = 100 and DisconnLeakRate = 0 for all service plans

Parameters for DirecPC™ Turbo-Throttling II

Data Sheet

The following data sheet explains the Turbo-Throttling parameters and lays down thumb rules for assigning values to these parameters. Turbo-Throttling maintains a leaky-bucket for each user; the following parameters are associated with the bucket (and are configurable on a service plan basis):

- L_1 - Leak rate (kbps): A user downloading data at this average rate would have a zero bucket depth (B , in bytes). The data sent to the user over the satellite causes increase in bucket depth. Every minute, a finite number of bytes is subtracted from the bucket depth to account for leaking.
- L_2 - Disconnected Leak Rate (kbps): The rate at which the bucket leaks when a user is not connected. This has been introduced to discourage users from keeping their modem connections up in order to reduce their running-average; especially if they use Bundled-ISP service plans.
- D - Duration (min): The duration over which the running average throughput (RAT) is calculated ($RAT = B/D$). RAT is displayed in Gateway statistics and is included in TIStats
- T_1 - Soft-Throttle Threshold (kbps): The average data rate that a user can maintain over a period of D minutes without encountering any throttling. A user is soft-throttled when his bucket depth becomes greater than $(T_1 - L_1) * D$ bytes
- T_2 - Hard-Throttle Threshold (kbps): The average data rate that a user can maintain over a period of D minutes without encountering any hard-throttling. A user is soft-throttled when his bucket depth becomes greater than $(T_2 - L_1) * D$ bytes
- P_0 - UnThrottled Peak Throughput (kbps): The maximum throughput that a user can get when not throttled. (The un-throttled window size = $P_0 * RTT$, RTT = Estimated Round-Trip Time)
- P_1 - Soft-Throttle Peak Throughput (kbps): The maximum throughput that a user can get when in soft-throttled state. (The soft-throttle window size = $P_1 * RTT$)
- P_2 - Hard-Throttle Peak Throughput (kbps): The maximum throughput that a user can get when in hard-throttled state. (The hard-throttle window size = $P_2 * RTT$)

The following spreadsheet provides sample values for MoonSurfer II:

DirecPC Turbo Throttling II Parameters (with sample settings for MoonsurferII)				
L_1	28 kbps	12.6 Mbyte/hr	{=Leak Rate*3600/8*1000}	MB/hr with 0 bucket depth
T_1	120 kbps	31.05 MByte	{=(T_1 - L_1)*D*60/8*1000}	MB in D min w/ no throttling
T_2	200 kbps	58.05 MByte	{=(T_2 - L_1)*D*60/8*1000}	MB in D min w/ no hard-throttling
D	45 min			
P_0	200 kbps	15000 bytes	{=P_0*RTT*1000/8}	Un-Throttled Window Size
P_1	112 kbps	8400 bytes	{=P_1*RTT*1000/8}	Soft-Throttle Window Size
P_2	56 kbps	4200 bytes	{=P_2*RTT*1000/8}	Hard-Throttle Window Size
RT	0.6 seconds			
Active Users on 24 Mbps Xponder @ Leak Rate (N)			771.43	{=24000*0.9/L_1}
% Active Users During Busy Hour (X)			8	
Users Per Transponder (U)			9642.9	{=N/(X/100)}
Revenue @ \$20 per user per month			192857	{=U*20}

The following thumb rules should be adhered to:

- T_1 , L_1 and D should be chosen such that a user can download at least 25 Mbytes (size of the biggest browser download) without any throttling.
- L_1 should equal economically sustainable average bit rate (ABR) for a service plan, that is, L_1 should be such that it satisfies the revenue requirements. In particular, the ABR for a service plan should be monitored and P_1 and P_2 should be adjusted until ABR equals L_1 (Note: To maintain a "happy" customer base, L_1 should be more than the ABR for enough users)
- P_0 equals the maximum data rate for the particular service plan; T_2 is typically $\leq P_0$.
- $P_1:P_2$ ratio should be approximately 2:1

Chapter 2 Overview

2.1 Overview

The Turbo-Throttling feature is designed to allow a fair access to NOC resources to all users and to provide performance-differentiated levels of service. It also allows throughput reduction (via window sizes) if average throughput over a time period exceeds (service-plan configurable) thresholds.

The feature also changes the definition of active minutes used in Billing records to equal the actual "connected" minutes (instead of the previous definition of "the minutes in which some user-data traverses the NOC").

The Billing system will be required to provide a flat-file which contains the site-id and service plan information for each user, say, on a daily basis.

2.2 Deliverables

- Updated Hybrid Gateway (WinNT) executables
- Updated DAK software
- SCO UNIX shell scripts and executables for uploading service-plan information from the Billing system and updating the Oracle Database
- Enhanced Oracle Forms to allow changes in Turbo-Throttling parameters (on a service-plan basis)
- Enhanced Oracle triggers to reconfigure Hybrid Gateways with change of service plan of individual users and/or change in Turbo-Throttling parameters
- ACS Upgrade Procedure
- Feature Design Document
- Feature Integration Test Plan
- Updated NOC Manual
- Updated Oracle Schema Document
- Updated DirecPC Product Technical Specification (description of ascii-file used for updating Service Plan information)

2.3 Compatibility and Expected System Impact

This section identifies the expected impact of this feature on various subsystems and specifies the compatibility requirements for the feature in tabular form (Table 4). The table indicates:

- Which system components are expected to be modified.
- Whether new NOC hardware or third-party software is required by the NOC for this feature. This is important as DirecPC franchises may be required to purchase the hardware or software.
- Whether the modified component is backwards compatible with the NOC, DAK or APPs as they existed prior to this feature. Backwards compatibility with earlier DAK and APP software is a requirement. Backwards compatibility with the NOC is highly desirable, especially for DAK and APP software as it allows the DAK and APP software to be released prior to upgrading the NOC.

Table 4. Compatibility and Expected System Impact

Chapter 3

Functional Requirements

3.1 Functional Requirements

The following paragraphs describe the basic functional requirements of the DirecPC Turbo-Throttling feature.

3.1.1 [R:0019] Running Average Throughput

The HGW will maintain a running average throughput for all users. The running average will be calculated by using leaky-bucket approach where the bucket size will depend on N, Running Average Duration (where N will be a configuration parameter with a default value of 60 minutes).

3.1.2 [R:0033] Assign initial window size based on Service Plan

Each user should be assigned an initial per-TCP-connection window size. This window size is equal to the WMax parameter for the corresponding Service Plan. The user's TCP window size should never exceed this value.

Starting with Turbo-Throttling II, each user will be assigned an initial Peak Throughput (PeakUnThrottled Throughput) which will be used, in turn, to calculate the initial TCP window size. (Window Size = Throughput * RTT, RTT = Round-Trip Time)

3.1.3 [D:0022] Exclude first few minutes of usage every day from running average calculation

In order to provide the user with a better initial experience, the HGW will not include the first X minutes (where X is a configuration parameter, with a default value of 0) of data transfer each day in the calculation of running average throughput for each user. During this time, the window size should be equal to the WMax parameter for the corresponding Service Plan (irrespective of the user's current window size). WMax >= 536.

3.1.4 [R:0034] Clamp window size based on running average throughput

If the running average throughput exceeds a (service-plan configurable) threshold T1, the user's TCP window size should be clamped to W1. Similarly, if the bucket size exceeds a higher threshold T2, the user's TCP window size should be clamped to W2. W1 and W2 will each be greater than 536 bytes and WMax >= W1 >= W2, T1 <= T2.

Starting with Turbo-Throttling II, the window sizes will be calculated based on the Peak Throughputs (which are configurable on a service-plan basis).

3.1.5 [R:0036] Per-user Throughput Throttling

When a user is throttled, the available bandwidth (or window size) will be divided among the "bulk-transfer" type connections. A bulk-transfer connection is a connection which, in the recent past, has had unacknowledged bytes exceeding a certain threshold (Other connections are termed as "interactive"). Hence all connections start as "interactive" and may become "bulk-transfer" later. These definitions correspond to the ones used for flow control.

3.1.6 [R:0035] Provide two levels of UDP discard corresponding to T1 and T2

Non-DNS UDP traffic should be discarded if the user is being throttled. Just like window size clamping, there should be 2 levels of UDP discard, U1 and U2 (corresponding to W1 and W2 respectively). For example, U1 = 50 implies that at the first level of throttling, every 50th UDP packet will be discarded. Note that 1 <= U2 <= U1 <= 1000.

3.1.7 [R:0037] TCP Connection Limiting

The number of TCP connections that can be opened simultaneously by a user or a site will be limited to the MaxConnections configured for the corresponding service plan. The statistics regarding Number of connections open at any instant will be available via Enhanced Turbo-Internet statistics on a real-time basis.

3.1.8 [R:0038] Bucket Leak

Turbo-Throttling I used the Bucket Leak Rate to be the same as the Throughput Hard Threshold (T2). Also, the bucket leak happened only if a user was active in a particular minute. Starting with Version 2 of Turbo-Throttling, new parameters will be provided to specify the Bucket Leak Rate when a user is connected and also when a user is not connected. The Disconnected Bucket Leak Rate is being added in order to discourage ISP Bundled users from keeping their phone line up (at HNS' expense) just to reduce their running average throughput.

3.1.9 [D:0023] Use “connected” minutes for active minutes and running average calculation

The HGW must use the actual “connected” minutes for calculating the “active” minutes and running average throughput. For this purpose, the DAK software must be modified to notify the HGW about connection establishment, connection tear-down and connection-up events (connection-up event must be notified at periodic intervals).

The HGW software must be modified to accept these messages and to time-out if connection tear-down message is not received for D minutes (where D is a configuration parameter, with a default value of 3 minutes). This must be independent of TCP connection timeouts (idle-timeout etc).

3.1.10 Persistent data at HGW

The following data at the HGW needs to be maintained across user sessions and across HGW restarts:

- [R:0025] The Last state of user (throttled/non-throttled), Number of minutes the user has been connected in the day
- [D:0026] The Current Bucket Depth for a user, Minutes the user has been under hysteresis and number of UDP packets transmitted without any discards

3.1.11 [R:0031] Disable data forwarding to a user when not logged on

By default, the HGW must not forward data to a user who is not currently logged on. Data forwarding should be allowed on a service-plan basis.

3.1.12 [R:0012] Update of Service Plan information for each user

A utility will be provided to update the Service Plan information for each user (in the ACS/Oracle database) with the data received from the Billing subsystem. The data will be in the form of ascii-file where each line of the file will have a 10-character site-id, a serial number and a 2-character numeric Service Plan-id, with comma as the separating character.

3.1.13 [R:0039] Unlimited number of service plans

The limit of 30 service plans (in version 1 of Turbo-Throttling) will now be removed.

3.1.14 [R:0013] HGW reconfiguration

The Hybrid Gateway reconfiguration must be triggered off whenever there is a change in the Turbo-Throttling parameters or a change in an individual user's service plan.

3.1.15 [R:0042] Enhanced TIStats

The following statistics will be added to the displayed statistics and also added to the hourly records:

- Number of TCP connections rejected
- Number of Hard Throttled minutes
- Maximum Bulk Transfer Connections in the last hour
- Bytes Received by a user in the last hour (already being displayed; added to hourly records)
- Service Plan of user (already being displayed; added to hourly records)
- Hybrid Gateway Id for the user (will just be added to hourly records)

3.2 Non-Throttling requirements

The following HGW enhancements will also be performed as part of Turbo-Throttling II:

3.2.1 [R:0040] Local Key Table

The HGW maintains a table to store the encryption keys for each remote. The keys are received from the CAC via multicast messages. These messages contain keys for all remotes. In the past, each HGW used to store all these keys. This behavior will be changed in the new release and only keys for remotes assigned to the particular HGW will be stored in the HGW.

3.2.2 [R:0041] Change in Flow Control (FC) mechanism

The current Flow Control mechanism provides for a FCTWS (Flow-Controlled Target Window Size). Each user's window size at any instant is the minimum of this Target Window Size and User's Window Size based on Throttled state. When a flow control message is received with latency greater than a threshold, the FCTWS is reduced by a certain percentage. This behavior was appropriate earlier; but with Turbo-Throttling, this leads to a condition wherein a hard-throttled user's window size may not be affected at all (since it was already less than the new FCTWS) while a non-throttled user's window size may start shrinking, ultimately becoming comparable to that of a hard-throttled user.

Hence, starting with Turbo-Throttling II, the HGW will maintain a Flow Control Meter (FCMeter) with values between 1 and 100. This FCMeter value will vary with the latency values reported by the SGW. The user's window size at any stage will be obtained by multiplying this value with the window size value based on his service plan and the state of throttling.

3.3 Obsolete Requirements

The following requirements were suggested in an earlier version of the Requirements document, but were subsequently made obsolete due to change in business plans.

3.3.1 [X:0020] Terrestrial redirection of user traffic

When the running average throughput of a user exceeds T bits per second (where T is a configuration parameter, with a default value of infinity == 10Mbps, typical value would be 56*1024 == 56kbps), the user traffic will be re-directed terrestrially.

3.3.2 [X:0021] Return to satellite path

Once a user has been terrestrially re-directed, the user will be switched back if both the following conditions are satisfied:

- The running average throughput falls below T bits per second
- The user's traffic has been terrestrial for at least M seconds (where M is a configuration parameter, with a default of $60*4 == 4$ minutes)

3.3.3 [X:0024] Updated Turbo-Internet Statistics

The Turbo-Internet Statistics must be enhanced to include the terrestrially-redirected bytes and the number of minutes of terrestrial redirection for each user in an hour. The running average throughput for each user must also be included in the Turbo-Internet Statistics.

3.3.4 [X:0027] Terrestrially redirect some portion of traffic at all times

The HGW must try to make use of the terrestrial bandwidth available by redirecting R kbps (where R is a configurable parameter, with a default value of 8kbps) of user traffic all the time.

3.3.5 [X:0028] Minimum session time

The NOC software should ensure that each user session lasts for at least S minutes (where S is a configurable parameter, with a default value of 5 minutes). This is required to prevent users from using the satellite only mode (in DAK) to download big chunks of data and then switch to Terrestrial only mode (in DAK). It should also be ensured that this implementation does not produce incorrect "connect minutes" in billing records.

3.3.6 [X:0029] Hourly usage notification to user

The HGW software will be enhanced to send a periodic usage-notification message to all configured DirecPC remotes (this activity should consume a very small satellite bandwidth). The DAK software will be enhanced to listen for these messages and update an hour-meter display which can be started by the user from within DirecPC Navigator Statistics.

The hour-meter must display actual usage v/s allotted quota per month and must pop up a dialog box when only 5 minutes of usage are left (or on session start with less than 5 minutes left).

3.3.7 [X:0030] Per-Service Plan, Per-HGW Turbo-Throttling parameters

The following configuration parameters will be configurable on a service-plan basis in the Auto-Commissioning Server. Based on the user's service plan and the corresponding service plan parameters, the HGW must use various Turbo-throttling mechanisms as described in the previous requirements.

- N, Number of minutes over which the running Average Throughput is calculated
- T, Threshold average Throughput for terrestrial redirection of user traffic
- M, Minimum number of seconds for which the user will be terrestrially redirected irrespective of whether the user's throughput falls below T before that period.
- X, eXcluded minutes: First few of minutes (on a daily basis) excluded from Running average Throughput calculation for each user
- R, Redirect-always throughput: Number of kbps of data which will be sent terrestrially all the time
- E, Enable Disconnect-Forward: If set, this will enable Data forwarding to a user who is not logged-on
- B, Down-Rev SW Block: If enabled, downrev software will not operate through the HGW. This would be employed for all users signing up for new service plans so that active minutes can be ensured to be the same as connected minutes.

3.3.8 New configuration parameters for Hybrid Gateway

The following new configuration parameters will be added (as new fields in the *hybridgws* table). All these parameters can be changed by the NOC operator and will immediately update Hybrid Gateway internal data structures after a Hybrid Gateway reconfiguration.

- [X:0001] PercentTerrRedirect: The percentage of users (for example, 10% of users in a list organized in descending order by throughput) in a particular priority level who will be moved to Terrestrial Redirection (simultaneously) on detection of congestion.
- [X:0002] CongestTimer: Whenever this timer goes off, the last received latency (from SGW) for each priority level is examined. If the value is greater than *LatencyThresh*, the *PercentTerrRedirect* highest-throughput users in that priority level will be switched to Terrestrial Redirection for duration equal to their *TerrRedirectMin* (*TerrRedirectMin* depends on Service Plan of user)
- User-specific configuration parameters: Each user's configuration record will be enhanced to include the following parameters (each of which is imported from the corresponding Turbo Internet service plan):
 - [X:0003] UserType: This determines the initial priority for a user in a particular service plan (that is, before any Throttling mechanisms are applied). This can have the following values:
 - 0 (default): "unlimited-access" user. The interactive traffic of such a user goes as Highest Priority while the bulk traffic will go as Medium, Low or Lowest Priority depending on the running average Throughput.
 - 1: "pay-per-byte" user. The interactive traffic of such a user goes as Highest Priority while the bulk traffic will go as High Priority irrespective of the running average Throughput.
 - [X:0004] TerrRedirectMin: Terrestrial Redirect Minutes (Default=5 minutes, value of 0 indicates no terrestrial redirection for this service plan).
 - [X:0005] ThruThreshMedPrio: Throughput Threshold for Medium Priority (default: 20kbps)
 - [X:0006] ThruThreshLowPrio: Throughput Threshold for Low Priority (default: 50kbps; If Throughput > ThruThreshLowPrio, user gets lowest priority)
 - [X:0007] RunningAvDuration: Running Average Duration used to calculate Throughput (default: 60 min, minimum: 10 min, maximum: 24*60 min)

3.3.9 New statistics at Hybrid Gateway and in Oracle Database

The Hybrid Gateways statistics gathering (and the Oracle database) will be enhanced to include the following new statistics:

- [X:0008] Mbytes downloaded/active minutes and terrestrial redirect bytes/minutes by user and priority level. These statistics should include total statistics and the per-application statistics and should be dumped on an hourly basis.
- [X:0009] Performance statistics which include: CPU usage, SGW latency values received (for each priority level), satellite throughput and terrestrial redirect throughput by priority level.

3.3.10 [X:0010] New priority levels at HGW and SGW

The following new priority levels (and associated Gateway Ids) will be introduced for TurboInternet traffic flowing from the Hybrid Gateway to the Satellite Gateway (Note: Priority 0 and 1 are reserved for Multimedia and Package Delivery traffic):

- Priority 2/Id 1: Interactive traffic (all service plans). This includes IP/UDP and light TCP traffic for all users

- Priority 3/Id 7: (High Priority) Bulk transfer traffic of per-Mbyte users (This includes the Basic and Bulk service plans)
- Priority 4/Id 8: (Medium Priority) Bulk transfer traffic of "unlimited access" (or non-pay-by-Mbyte) users whose average throughput is less than *ThruThreshMedPrio*
- Priority 5/Id 9: (Low Priority) Bulk transfer traffic of "unlimited access" (or non-pay-by-Mbyte) users whose average throughput is greater than *ThruThreshMedPrio* but less than *ThruThreshLowPrio*
- Priority 6/Id 10: (Lowest Priority) Bulk transfer traffic of "unlimited access" (or non-pay-by-Mbyte) users whose average throughput is greater than *ThruThreshLowPrio*

3.3.11 [X:0011] Terrestrial return of Turbo-Internet traffic

The hybrid gateway should allow for terrestrial redirection of all traffic for a user for a configured duration. Initially, only bulk TCP traffic will be terrestrially redirected.

3.3.12 Hour-meter display

The DAK software will provide a display of the hourly usage over the current month and the allotted hours for the month. This is covered under [X:0029]

3.4 [R:0032] User Interface

All user interface changes must be approved by the lead engineer responsible for user-interface design and development. All internationalization issues involved in such changes must also be considered.

3.4.1 [R:0014] Configuration of Turbo-Throttling parameters via Oracle Forms

The Oracle Forms must be enhanced to allow configuration of Turbo-Throttling parameters (mentioned in [R:0030]) by the NOC operator. This enhancement should take into account any issues related to the internationalization features (e.g., "internationalized" character set) inherent in Oracle 7.3 and later releases.

3.5 External Interface Requirements

This section intentionally left blank.

3.5.1 User Interfaces/Characteristics

This section intentionally left blank.

3.5.2 Hardware Interfaces

This section intentionally left blank.

3.5.3 Software Interfaces

This section intentionally left blank.

3.5.4 Communications Interfaces

This section intentionally left blank.

3.5.5 Network Management Interfaces

This section intentionally left blank.

3.6 Expected Enhancements

This section intentionally left blank.

Chapter 4

System Attributes

This section documents the non-functional requirements of DirecPC Turbo Throttling feature.

4.1 Performance

4.1.1 [R:0015] No impact on existing Billing/Performance statistics

Turbo Throttling must not affect the Billing/Performance statistics being gathered at present, except that the active minutes will now equal "connected" minutes.

4.2 Capacity and Resources

This section intentionally left blank.

4.3 Scalability

This section intentionally left blank.

4.4 Design Constraints

This section intentionally left blank.

4.4.1 [R:0016] Standards Compliance

All phases of implementing the requirements set forth in this document shall be ISO 9001 compliant as set forth by HNS ISO 9001 procedures and practices.

4.4.2 Hardware Limitations

This section intentionally left blank.

4.4.3 Physical Environment Considerations

This section intentionally left blank.

4.4.4 Timing Constraints

This section intentionally left blank.

4.4.5 [R:0017] Software Compatibility/Environment Considerations

The feature must operate on Windows NT HGW and must be compatible with other OS/2 or WinNT MUX components.

4.4.5.1 [R:0043] Backwards Compatibility

1. If the database is upgraded without a particular HGW being upgraded, that HGW will just ignore the new parameters. Or if the new version of HGW is replaced by the older version, it will continue to function as before.
2. If the HGW is upgraded without a database upgrade, the HGW will assume a value of 100 for MaxTCPConnections for all service plans. Also it will assume Throttling I, with DisconnLeakRate = 0.

4.4.6 Unit Cost

This section intentionally left blank.

4.5 Time to Market

This section intentionally left blank.

4.6 Testability

This section intentionally left blank.

4.7 Attributes

This section intentionally left blank.

4.7.1 Availability

This section intentionally left blank.

4.7.1.1 Startup/Recover Time

This section intentionally left blank.

4.7.1.2 Fault Tolerance

This section intentionally left blank.

4.7.1.3 Mean Time Between Failure (MTBF)

This section intentionally left blank.

4.7.2 Security

This section intentionally left blank.

4.7.3 Privacy

This section intentionally left blank.

4.7.4 Debugging/Trace Capability

This section intentionally left blank.

4.7.5 [R.0018] Maintainability

The following new documents will be created for this feature:

- DirecPC Turbo Throttling Feature Requirements Specification (Pubs No. HNS-11232)
- DirecPC Turbo Throttling Feature Design Specification (Pubs No. HNS-11422)
- DirecPC Turbo Throttling Feature Integration Test Plan (Pubs No. HNS-12090)

The following documents will be updated for this feature:

- DirecPC NOC Manual (Pubs No. HNS-7013).
- DirecPC Hybrid Internet Technical Specification (Pubs No. 8050744)
- DirecPC Oracle Schema Document (Pubs No. HNS-6896)
- DirecPC Product Technical Specification (Pubs No. 8050134)

The following are the other deliverables:

- Executables for WinNT-based HGW
- New DAK software
- SCO UNIX shell scripts and executables for uploading service-plan information from the Billing system and updating the Oracle Database
- Enhanced Oracle Forms to allow changes in Turbo-Throttling parameters (on a service-plan basis)
- Enhanced Oracle triggers to reconfigure Hybrid Gateways with change of service plan of individual users and/or change in Turbo-Throttling parameters
- ACS Upgrade Procedure

4.7.6 Adaptability

This section intentionally left blank.

4.7.7 Portability

This section intentionally left blank.

4.8 Installation Requirements

This section intentionally left blank.

4.9 Customization Requirements

This section intentionally left blank.

4.10 Other Requirements

This section intentionally left blank.